



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) Publication number:

0 529 130 A1

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: 91114533.2

(51) Int. Cl. 5: G06F 9/44

(22) Date of filing: 28.08.91

(43) Date of publication of application:
03.03.93 Bulletin 93/09

(84) Designated Contracting States:
DE FR GB

(71) Applicant: International Business Machines Corporation
Old Orchard Road
Armonk, N.Y. 10504(US)

(72) Inventor: Maler, Franziska
Schwarzwaldstrasse 324

W-7800 Freiburg(DE)
Inventor: Müller, Claudia
Sonnenmatte 9
W-7620 Oberwolfach(DE)
Inventor: Saelens, Serge
Steinstrasse 32
W-7044 Ehningen(DE)

(74) Representative: Jost, Ottokar, Dipl.-Ing.
IBM Deutschland GmbH Patentwesen und
Urheberrecht Schönaicher Strasse 220
W-7030 Böblingen (DE)

(54) Expert system with explanation system.

(57) An explanation system for use in an expert system is described. The rule base used in the inference engine (20) of the expert system is divided into groups of rules (50) called rule classes. With each rule class, three types of explanations (60) are associated: strategy explanations, reason explanations and inference explanations. The rules classes are arranged in a hierarchical explanation tree structure. Each time a first rule is fired within rule class, a frame is created into which the explanations relating to that rule class are copied (340, 360). After the completion of the reasoning process, the frames can be used to generate an explanation of the reasoning process.

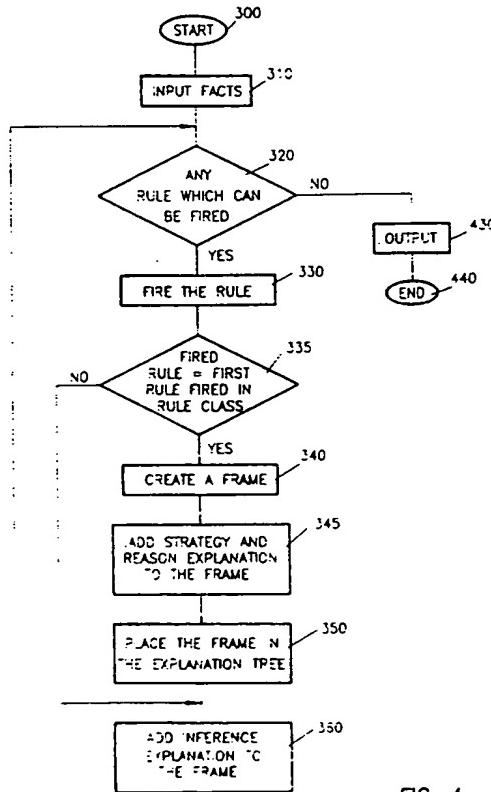


FIG. 4

EP 0 529 130 A1

Field of the Invention

The invention concerns an expert system with an inference engine, a rule base containing rules for the operation of the expert system and an explanation system for producing an explanation of the operation of the expert system.

Prior Art

Expert systems or knowledge-based systems are being increasingly used to help in a variety of fields, such as the diagnosis of illness or identification of faults in a machine. They use whatever facts are available to either produce a result or confirm a suspected result. In carrying out the reasoning a number of pre-programmed rules, called a rule base, are used which are developed from the knowledge of experts in the field in which the expert system operates. These rules are processed by what is usually termed an 'inference engine'.

A number of different tools are known for allowing the development of an expert system. The PROLOG Language and IBM Knowledge Tool are two well known of these tools. US-A-4 803 641 (Hardy et al) also describes a tool for developing an expert system.

Two types of mechanism are most usually employed in the inference engine in order to carry out the reasoning process. An inference engine with a forward chaining procedure starts with an initial set of facts and develops from this initial set as many further facts as possible. An inference engine with a backward chaining procedure on the other hand is goal-oriented. It starts with a premise (goal) which must be given at the beginning of the procedure and then tries to see whether the given facts confirm this goal. In European Patent Application EP91105698.4, an inference engine is described in which the forward-chaining and backward-chaining methods are integrated.

In addition to knowing the result of the analysis, the user is often interested in knowing how and why the inference engine came to the result that it did. This requires the presence of an explanation system to produce an explanation of the reasoning that was carried out. The user may also be interested in hypothetical possibilities, such as what would happen if certain other facts were given or he may be interested in knowing why the inference engine did not come to a certain conclusion. Such questions might also be answered in an explanation system. A survey of explanation systems and possible approaches is given in the article "Explanation in intelligent systems" by C. Ellis which was published in "Expert Knowledge and Explanation: the knowledge-language interface" edited by C. Ellis, Ellis Horwood, Chichester, 1989.

In the simplest expert systems, explanation systems merely involved a listing of all the rules that were 'fired' (i.e. used) during the analysis. An example of such a system is the MYCIN system which was developed to assist a physician who was not an expert in antibiotics with the diagnosis and treatment of bacterial blood infections. The MYCIN system is described in the book "Rule-based expert systems: the MYCIN Experiments of the Stanford Programming Project", edited by B.G. Buchanan and E.H. Shortliffe, Addison-Wesley, Reading, Massachusetts, 1984.

However, as expert systems become increasingly complicated, such a procedure is inefficient since many hundreds of rules may be fired during the course of the reasoning process. In addition the comprehensibility of the explanation for the average user decreases as he has to search among the large number of rules displayed in order to find the explanation that he is searching for. Finally, since complete explanations are stored for every rule fired, the performance of the system is diminished since a number of duplicate explanations may be stored.

A further problem encountered when using explanation systems in an expert system was identified by C. Millet and M. Gilloux in their article "A Study of the Knowledge Required for Explanation in Expert Systems" published in the "Proceedings of the Fifth Conference on Artificial Intelligence Applications (CAIA)", 6-10 March 1989, Miami, Florida, pp 83-90 and published by the IEEE Computer Society Press. They realised that in existing implementations of expert systems part of the knowledge that should be transferred was not communicated. They identified two types of uncommunicated knowledge: 'missing knowledge' and 'implicit knowledge'. Missing knowledge is that knowledge which was used by the expert in designing the rules for the expert system which is, however, not required during the search for a solution. Implicit knowledge is the knowledge of which the expert may not be aware that he possesses but nonetheless he 'implicitly' assumes when he is designing the rules for the expert system. An example of such implicit knowledge is the way in which the rule is expressed. Whilst the wording may be clear for the expert, the average user of the system may fail to appreciate the significance of expressing the rule in one way rather than in another way. Their solution to this problem was to build another knowledge base on top of the existing rule base which did not change the reasoning approach of the system but was only limited to explanatory purposes. The inclusion of this extra knowledge base reduces the performance of the system substantially since, in addition to carrying out the reasoning process, the expert system has

to fetch from the extra knowledge base the required explanations.

R.Rubinoff in his article "Explaining Concepts in Expert Systems: the CLEAR System" published in "Artificial Intelligence Applications: The Engineering of Knowledge-based Systems: Proceedings of the Second Conference, Miami Beach, Florida, December 11-13, 1985", pp 416-421, edited by C.R.Weisbin and published by the IEEE CS Press: North Holland, 1985 identified a further need that users of expert systems required when they were confused about what the system was asking. He suggested 'generalised' rules be added to the rule base. These rules themselves could be further generalised thus creating a hierarchy of rules. Thus, whenever several rules for explaining a concept are selected that the user will not understand and these rules are all grouped together in a generalised rule, then the generalised rule will be displayed instead of the several individual rules. This type of explanation system, however, requires the developer to develop increasingly abstract rules which may or may not be helpful to the understanding of the reasoning process.

Further types of explanation systems are known. However, they often require special adaptations of the rule base to incorporate the explanation system. This adaptation alters the structure of the rule base and, as a result, can not be easily incorporated into existing rule bases.

Summary of the Invention

The object of the invention is to build an explanation system for incorporation into an expert system which overcomes the above problems.

A further object of this invention is to build an explanation system which can be incorporated into an existing expert system without requiring changes in the structure of the rule base.

This is achieved by grouping the rules in the rule base of the expert system into rule classes. To each rule class, a number of types of explanations are added, such as strategy explanations or reason explanations. In addition, further explanations can be added as the rules are fired.

The explanations are collected together in a series of frames and, after the reasoning process has been completed, the user can interrogate the frames in order to understand the reasoning process.

The user can use a number of methods in order to interrogate the frames in order to obtain explanations about the reasoning process. For example windowing techniques can be used to allow the user to run through the explanations on a graphics screen using a mouse. Alternatively a command language could be developed to allow

the user to control the manner in which explanations are given.

The explanation system as described finds applications in all fields in which expert systems are implemented. It can be used in medical systems to allow paramedics to safely diagnose and treat emergencies or in allowing a computer repairman to understand the output of a computer dump to analyse a failure in a computer system.

Figures

- Fig. 1 shows the structure of the invention.
- Fig. 2 shows the structure of a rule class.
- Fig. 3 shows the structure of an explanation tree.
- Fig. 4 shows a flow diagram for the generation of explanation frames during process.

Detailed Description

The structure of the apparatus used in this invention is depicted in Fig.1. It comprises an input 10 in which the known facts are input, an inference engine 20 in which the facts are analysed and inferences drawn from the facts, a rule base 25 containing rules used to analyse the facts and an explanation system 30 which provides explanations about the reasoning process. The inference engine 20 can use any type of inference mechanism; in the preferred embodiment a forward-chaining mechanism is used. The rule base 25 and the explanation system 30 are both connected to an output 40. The output 40 can be any type of output device capable of displaying text such as a printer or a display screen.

The rules forming the rule base 25 and processed by the inference engine 20 are grouped together in groups which are called rule classes. The data structure of one such rule class is shown in Fig. 2. Each rule class comprises a series of rules which all contribute to the same goal in the reasoning process. These rules are stored in a rule table 50. For example, in an expert system for diagnosing illness, one rule class may contain rules which test whether the illness is caused by a bacteria, whilst another rule class may contain rules to determine whether a virus is involved. The rules may be of the form IF....THEN or any other form. The rules 50 do not need to be grouped together in the rule base 25. All that is necessary is that each rule be assigned a rule class to which it belongs.

With each rule class, explanations can be associated. In the example shown in Fig. 2, three types of explanations are provided. Two of these types of explanations are given in an explanation table 60.

Strategy explanations are the general principles according to which the reasoning steps are carried

out. For each rule class, the strategy explanations will indicate the function of the rule class, i.e. what is happening. An example of a strategy explanation is "diagnose illness", or "analyse problem".

Reason explanations give the background knowledge which forms the basis for the reasoning process. They give explanations about why something happens. Reason explanations include "Block A contains important information about the problem, therefore it makes sense to look there for background information". Inference explanations on the other hand are not taken from the explanation table 60 but are derived from the rule base itself as outlined below. This type of explanation may vary during every reasoning process as they comprise concrete statements which are derived during the reasoning process from the input facts. Inference explanations inform the user about how the function of the rule class had been fulfilled and what results have been achieved. An example of an inference explanation would be "as the patient has a temperature, runny nose and a headache, it is concluded that the patient has influenza".

The rule classes can be related together in a tree structure as shown in Fig. 3. Each of the rule classes form a node 100-270 of the tree. The nodes 100-270 of the tree can be connected together either through AND operations or through OR operations. The position of the rule class in the tree and its relationship to other rule classes in the same level of the tree is indicated in the dependence table 65 shown in Fig. 2.

This dependence table 65 shows which rule classes in a lower level of the tree are dependent on rule classes in the next higher level of the tree and also indicates which rule classes in a lower level of the tree are dependent on other rule classes in the same lower level. This means that if any of the rules in the rule table 50, for example that of the rule classes at nodes 130 and 140 are fired, then their generated explanations have to be read in conjunction with the explanations generated by firing rules in the rule table 50 of the rule class at node 110. For reasons of clarity, a simple tree structure is shown in Fig. 3. However, it is conceivable to develop more complicated structures in which nodes in lower levels of the trees are not dependent on only one node in a higher level, but on several nodes. For example node 130 in Fig. 3 is shown to be only dependent on node 110. However, it might also be made to be dependent on node 120. This relationship would also be indicated in the dependence table 65, thus forming an explanation net.

AND operations (such as those shown on Fig. 3 at nodes 120, 130, 150 and 160) connect nodes in the same tree level which depend on each other. An example of such a dependence would be the

relationship where a first node searches for a problem and the second node names the problem. Clearly reading the explanations generated by the second node is only necessary if interrogation of the rule class in the first node has shown that a problem exists. The order in which the rule classes are to be tested is also given in the dependence table 65 of the nodes.

OR operations connect nodes which do not depend on each other (such as those shown on Fig. 3 at nodes 100, 110, 130, 140 and 160). For example a first node might be "buy bread", a second node might be "buy sausage" and a third node might be "buy milk". Each of these operations is independent of each other, they can be carried out separately from each other and the generated explanations read separately from each other. The act of carrying out one operation does not affect the outcome of the other operations.

It can be seen at nodes 130 and 160 in the tree that some of the rule classes dependent on this node in the tree are connected by AND operations (nodes 170, 180; 230, 240) whilst others dependent on the same node are connected by OR operations (nodes 190; 240, 250, 260, 270).

The root 100 of the explanation tree give the most abstract explanations whilst the leaves 110-270 of the tree give the most detailed explanations. That is to say, the farther that one travels up the tree, the more detailed the explanations become. These detailed explanations allow the user of the system to obtain the required information about any small bit of the system that he or she requires.

Using the tree, the complete reasoning strategy followed by the expert system in coming to a conclusion from the input facts can be understood. This is possible since every time a rule is fired in one of the rule classes a strategy explanation will be produced explaining what the rule class is testing for. Thus by analysing all of the strategy explanations, the user can determine the procedure which the expert system used in coming to its conclusions.

The operation of the expert system will now be described. Before the expert system with the explanation system can be operated, one or more experts have to construct the rule base which is to be used. In addition to defining the rules, the experts have to classify the rules into rule classes. As mentioned above these rule classes group the rules together which fulfill the same task in the reasoning process.

With each rule class, the experts have to supply explanations associated with the rule class. In the preferred embodiment these explanations comprise strategy explanations or reason explanations. However, it is possible that other types of explanations can be provided. These explanations should

be of sufficient length and clarity to allow the user to understand the reasoning of the system. It is, of course, conceivable that the background knowledge of the users using the expert system will vary. Therefore it would be possible to grade the extent of the explanation. For example a novice user would require a much more detailed degree of explanation than a more experienced user. The different grade of explanations would be stored in the table 60 shown in Fig. 2 and the user of the expert system would then select which grade of explanation she or he requires. In addition a hypertext system could be installed to further explain unfamiliar terms.

As the expert system is used and as more knowledge about the problem is acquired, it is possible to upgrade the rule base of the expert system. This is simply done by the expert who upgrades the system selecting the rule class into which the new rule is to be placed and then adding a new rule to this class. For example in the example shown in Fig. 2, the expert would add a new rule n+1 to the rule base 50. The expert is not required to alter the strategy, reason or other explanations associated with each rule class.

Further rule classes can be additionally created without difficulty. This process would be carried out by the expert defining the new rules or reclassifying existing rules that are required and putting these into a new rule table 50, writing explanations about the function of the rule class and adding these to the explanation table 60 and finally defining the dependence of the rule class and storing this in the dependence table 65.

After the expert has defined the rules and added the explanations, the expert system is ready for use by the user. In Fig. 4 a flow diagram is presented showing the various operations that are carried out. In step 310 of the table the user inputs the facts that he knows and from which he wishes conclusions to be drawn. The expert system then searches the rule base 25 (step 320) to see whether any of these facts fire any of the rules in the rule base 25. This rule will then be fired (step 330). The fired rule does not necessarily need to be in the rule class forming the root of the tree, since the tree does not determine the order of rule firing. The tree is used for generating the explanations.

Should the first rule in any rule class be fired (step 335), then at step 340 a frame is created into which is placed the strategy explanation and the reason explanation for the rule class in which the fired rule is present (step 345). Using the information stored in the dependence table 65, the created frame can be placed automatically at the correct node in an explanation tree (step 350). In step 360 the inference explanation is also added to the frame, this type of explanation being directly de-

rived from the rule fired.

If, however, in step 335, it was determined that the fired rule was not the first rule in the rule class to have been fired, then no new frame needs to be created. Thus the process jumps automatically to step 360 and a new inference explanation is added to the frame. Since the strategy explanation and reason explanation for the rule class have already been included in the frame, it is no longer necessary to add them.

After the inference explanation has been added (step 360), the reasoning process returns to step 320 to see whether any more rules remain in the rule base to be fired. If this is the case then the reasoning process continues through the loop. If, however, no more rules remain in the rule base 25 to be fired, then an output is prepared in step 430.

The output in step 430 is constructed from the frames that have been created during the reasoning process and the explanation tree in which the frame has been placed. Each frame contains a strategy explanation, a reason explanation and a number of inference explanations depending on the rules fired in each rule class.

The level and quality of the output depends on which frames were created during the reasoning process and placed in the explanation tree. For example should rules have been fired which are in the leaves of the explanation tree, then fairly detailed explanations will be supplied. On the other hand should only rules classes in nodes deeper in the tree have been fired, the explanations will be much more general.

Using the frames, the user can then look through the supplied explanations with help of the dependence tables 65 to follow the reasoning strategy used and understand why certain conclusions were reached. A number of different support programs can be envisaged to allow the user to look at the frames. For example a graphics interface using windows technology would allow the explanation tree to be reconstructed. Alternatively a command language might be used to allow the user to directly proceed to certain points within the explanation. For a novice user who is not interested in the detailed explanations, an interface could be developed in which the only explanation given was a command such as "give penicillin urgently as suffering from illness" or "replace card No. 2378 as it is not functioning".

The explanation system described herewithin can be applied to any type of expert system in which explanations are required. It can be added to already existing systems developed using languages such as TIRS, PROLOG or the IBM KnowledgeTool system with a minimum of extra programming and overhead. This is possible since the structure of the rule base 25 does not need to

be altered. All that needs to be done is that each rule in the rule base 25 be associated with a rule class.

The explanation system will find wide applications in field such as analysing the memory dump from a computer system to determine why a program failed or in medical systems to allow paramedics to accurately diagnose problems and give quickly help.

Claims

1. Expert system with an inference engine (20), a rule base (25) for the operation of the expert system and an explanation system (30) for producing an explanation of the operation of the expert system

characterised in that

rules (50) in the rule base (25) are grouped into rule classes, whereby with each rule class explanations (60) are associated;

2. Expert system according to claim 1 further characterised in that

said explanations (60) include strategy explanations, inference explanations or reason explanations.

3. Expert system according to any of the above claims characterised in that

on firing a rule in a rule class, the explanations are stored in an explanation frame, which can be accessed by said explanation system (30).

4. Expert system according to any of the above claims characterised in that

a dependence table (65) is associated with each rule class to indicate the dependence of the rule classes on each other.

5. Expert system according to any of the above claims further characterised in that

said inference engine (20) operates in a forward chaining process, a backward chaining process or a mixture of the two.

6. Method for producing explanations in an expert system comprising the following steps

producing a set of rules (50) representing the knowledge about a system and storing them in an rule base (25);

5

providing (310) facts to an inference engine (20);

10

processing (320, 330) the facts in an inference engine (20) using the rules in the rule base (25);

15

producing (340, 345, 350, 360) explanations in the explanation system (30);

20

providing (430) conclusions from the inference engine (20) and explanations from the explanation system (30)

25

characterised in that

20

the step of producing a set of rules further comprises a step of grouping the rules together in a rule class and

25

associating with each rule class explanations (60).

30

7. Method according to claim 6 characterised in that

30

the step of grouping the rules together in a rule class further comprises defining a dependence table (65) to be associated with rule class and which shows the dependence of one rule class on another.

35

8. Method according to one of claims 6 or 7 characterised in that

35

the step of producing explanations in the explanation system (30) comprises creating (340) a frame and adding the explanations to said frame (345, 360).

40

9. Method according to claim 8 characterised in that

40

the step of providing (430) explanations from the explanation system (30) comprises interrogating the frame and producing a display on a display device.

45

10. Method according to any of claims 6 to 9 characterised in that

45

50

the step of processing (330) the facts in an inference engine (20) is carried out using a forward chaining process, a backward chaining process or a mixture of the two.

55

11. Use of the apparatus according to claim 1 to 6
for producing explanations about the problems
encountered in analysing a computer memory
dump.

5

10

15

20

25

30

35

40

45

50

55

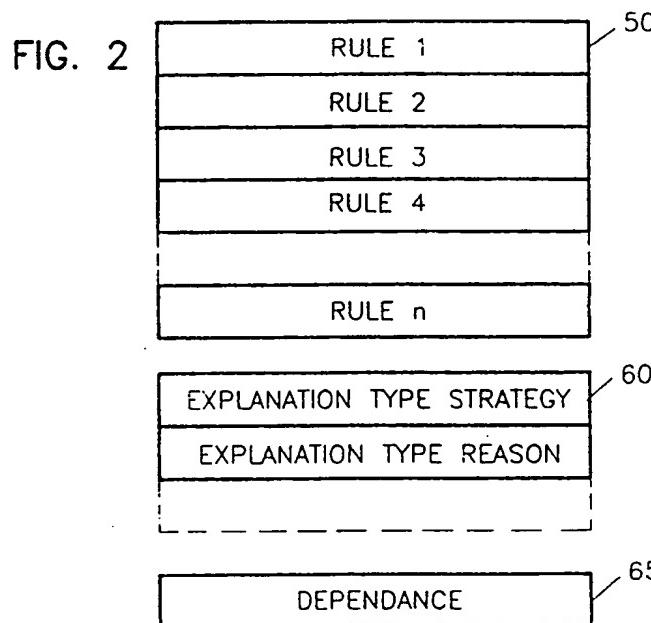
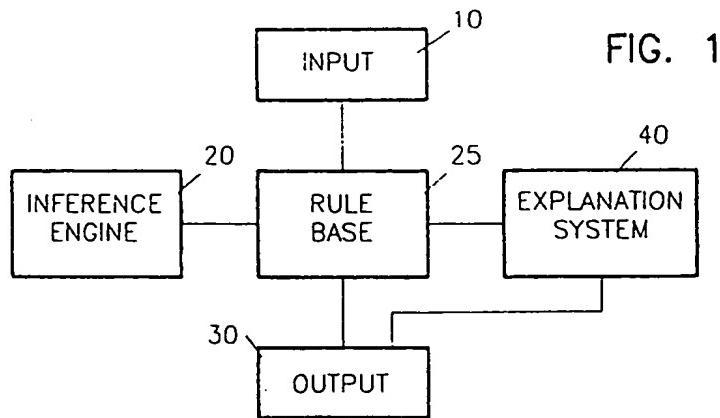
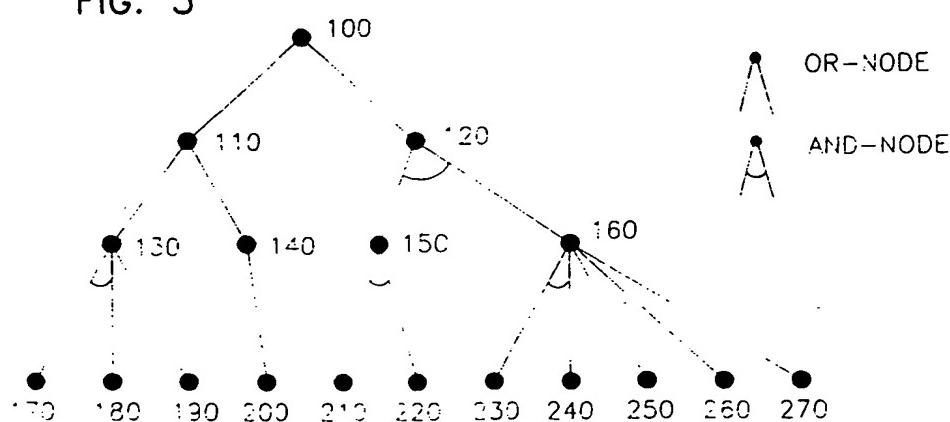


FIG. 3



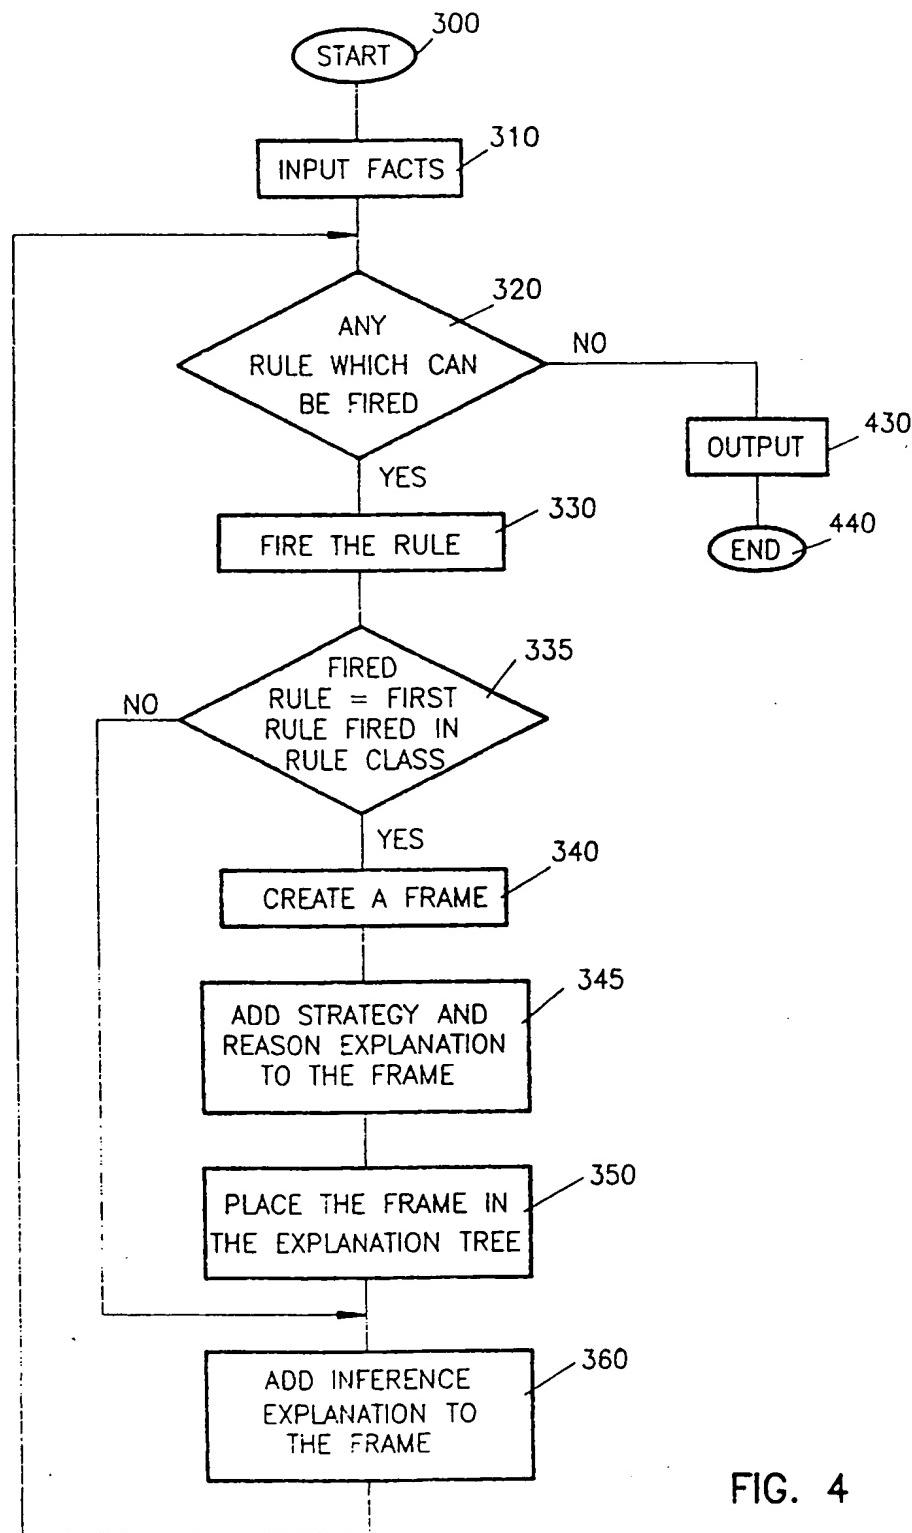


FIG. 4



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 91 11 4533
Page 1

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl.5)
D,X	<p>PROC. OF THE FIFTH CONF. ON ARTIFICIAL INTELLIGENCE APPLICATIONS 6 March 1989, IEEE, WASHINGTON, US pages 83 - 90; C. MILLET, M. GILLOUX: 'A Study of the Knowledge Required for Explanation in Expert Systems' * abstract * * page 85, left column, line 33 - line 34 * * page 86, left column, line 3 - line 9 * * page 88, left column, line 41 - page 89, left column, line 24 *</p>	1,2,5,6, 10	G06F9/44
D,Y	---	3,8	
D,X	<p>PROC. OF THE SECOND CONF. ON ARTIFICIAL INTELLIGENCE APPLICATIONS: THE ENGINEERING OF KNOWLEDGE BASED SYSTEMS 11 December 1985, IEEE, WASHINGTON, US pages 416 - 421; R. RUBINOFF: 'Explaining Concepts in Expert Systems' * abstract * * page 420, right column, line 25 - line 41 *</p>	1,2,5,6, 10	
D,A	<p>C. ELLIS 'Expert Knowledge and Explanation' 1989, ELLIS HORWOOD LTD., CHICHESTER, ENGLAND Ch. 5: "Explanation in intelligent systems", p108-126 * page 122, line 32 - line 41 *</p>	3,8	G06F
Y	<p>23RD ANNUAL MEETING OF THE ASSOCIATION FOR COMPUTATIONAL LINGUISTICS, ASSOC. OF COMPUTER LINGUISTICS 8 July 1987, MORRISTOWN, NJ, US pages 228 - 237; KUKICH, K.: 'Explanation Structures in XSEL' * page 228, left column, line 1 - line 20 * * page 231, right column, line 35 - page 232, left column, line 19 * * page 232, right column, line 11 - page 233, right column, line 10 *</p>	3,8	
<p>The present search report has been drawn up for all claims</p>			
Place of search	Date of completion of the search	Examiner	
THE HAGUE	16 APRIL 1992	KINGMA Y.	
CATEGORY OF CITED DOCUMENTS		<p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons A : member of the same patent family, corresponding document</p>	
<p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p>			



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 91 11 4533
Page 2

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. CL.5)
A	<p>ARTIFICIAL INTELLIGENCE vol. 15, no. 1-2, 1980, NORTH-HOLLAND, AMSTERDAM, NETHERLANDS pages 19 - 48; J.L. WEINER: 'BLAH, A System Which Explains its Reasoning' * abstract * * page 23, line 20 - line 31; figure 1 * * page 25, line 32 - page 25, line 6 *</p> <p>-----</p>	3	<p>-----</p> <p>TECHNICAL FIELDS SEARCHED (Int. CL.5)</p> <p>-----</p>
<p>The present search report has been drawn up for all claims</p>			
Place of search	Date of compilation of the search	Examiner	
THE HAGUE	16 APRIL 1992	KINGMA Y.	
CATEGORY OF CITED DOCUMENTS		<p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons A : member of the same patent family, corresponding document O : non-written disclosure P : intermediate document</p>	
<p>EPO FORM 1501 (02/82) (POB1)</p>			